

A Generic Online Parallel Learning Framework for Large Margin Models

Shuming Ma and Xu Sun

MOE Key Laboratory of Computational Linguistics, Peking University
School of Electronics Engineering and Computer Science, Peking University
{shumingma, xusun}@pku.edu.cn

Abstract

To speed up the training process, many existing systems use parallel technology for online learning algorithms. However, most research mainly focus on stochastic gradient descent (SGD) instead of other algorithms. We propose a generic online parallel learning framework for large margin models, and also analyze our framework on popular large margin algorithms, including MIRA and Structured Perceptron. Our framework is lock-free and easy to implement on existing systems. Experiments show that systems with our framework can gain near linear speed up by increasing running threads, and with no loss in accuracy.

1 Introduction

Large margin models have been widely used in natural language processing for faster learning rate and smaller computational cost. However, the algorithms may still suffer from slow training time when training examples are extremely massive, the weight vector is large, or the inference process is slow. With parallel algorithms, we can make better use of our multi-core machine and reduce the time cost of training process.

Unluckily, most studies about parallel algorithms mainly focus on SGD. Recht et.al (2011) first proposed a lock-free parallel SGD algorithm called HOGWILD. It is a simple and effective algorithm which outperforms non-parallel algorithms by an order of magnitude. Lian et.al (2015) provide theoretical analysis of asynchronous parallel SGD for nonconvex optimization and a more precise description for lock-free implementation on shared memory system.

Zinkevich et.al (2010) proposed a parallel algorithm for multi-machine called Parallel Stochastic Gradient Descent (PSGD). PSGD is an effective parallel approach on distributed machine but Recht et.al (2011) found that it is not as promising as HOGWILD on a single machine. Zhao and Li (2016) propose a fast asynchronous parallel SGD approach with convergence guarantee. The method has a much faster convergence rate than HOGWILD.

To the best of our knowledge, there is no related research about asynchronous parallel method for large margin models. McDonald et.al (2010) proposed a distributed structured perceptron algorithm but it needs multi-machine. We first propose a generic online parallel learning framework for large margin models. In our framework, each thread updates the weight vector independently without any extra operation or lock. Besides, we analyze the performance of structured perceptron and Margin Infused Relaxed Algorithm (MIRA) in our framework. The contribution of our framework can be outlined as follow:

- Our framework is generic and suitable for most of the large margin algorithms. It is simple and can be easily implemented on existing systems with large margin models.
- The framework does not use extra memory. Experiments show that the memory cost is no more than single-thread algorithm. Besides, the parallel framework works on a shared memory system so we have no need to care about the data exchange.

2 Generic Parallel Learning Framework

Suppose we have a training dataset with N samples denoted as $\{(x_i, y_i)\}_{i=1}^N$, where x_i is a sequence (usually a sentence in natural language

Algorithm 1 The Generic Parallel Framework

Input: training set S with N samples

- 1: **initialize:** weight vector $w = 0, v = 0$
- 2: **for** $t = 1$ to T **do**
- 3: Random shuffle training set S
- 4: Split training set into K part $\{S_i\}_{i=1}^K$
- 5: **for all** threads parallel **do**
- 6: Inference for update term Φ_i
- 7: Update w_{i+1} with Φ_i
- 8: $v = v + w_{i+1}$
- 9: $i = i + 1$
- 10: **end for**
- 11: **end for**

Output: the learned weights $w^* = v/(N * T)$

processing) and y_i is a structure on sequence x_i (usually a tag list or a tree). The whole dataset is trained with T passes. We denote the weight vector as w and after i^{th} update the weight vector will be w^i .

In each pass, the online learning algorithm will shuffle the dataset after which the weight vector is updated with each sample in the dataset. Generally, gradient-based algorithms like SGD take a lot of time to compute the gradient. Large margin algorithms like perceptron and MIRA spend most of time in decoding. Popular approach to speed up inference process is to take approximate inference instead of exact inference (Huang et al., 2012). However, we can parallel the inference process of several samples and update asynchronously to further accelerate the training process (Sun, 2016).

In our parallel framework, we split the dataset into k parts and then assign these split datasets to k threads. Each thread updates independently with a shared memory system. After that, we average the weight vector by the number of iterations (not the number of threads as some distributed parallel algorithms). We can see that our approach has no more computation than large margin algorithms, so it is a simple parallel framework. Since the whole framework runs on a shared memory system, there are mainly two problems about this framework: First, several threads update at the same time so it may be closer to minibatch algorithm instead of online learning algorithm intuitively. Whether the parallel framework will affect the convergence rate of online learning algorithm is a problem. Second, when a thread is working, the weight vector may be overwritten by other threads. Whether it will lead to divergence is an

other problem need to be analyzed.

For the first problem, Lian et.al (2015) proved that the convergence rate will not be affected under the parallel SGD framework. Our experiments also support that the convergence rate of large margin algorithms is still the same in our parallel framework. For the second problem, Recht et.al (2011) shows that individual SGD steps only modify a small part of the decision variable so memory overwrites are rare and barely any error will be made into the computation. We will also explain this problem on large margin models in Section 3. Experiments show that our parallel algorithm is so robust that the interference among threads will not affect the convergence. In our framework we also average the weight vector by the number of iterations. One reason is that Collins (2002) explains that averaging parameter helps avoid overfitting. Another advantage is that we can ensure every update will contribute to the final learned weight vector.

3 Large Margin Models

In this section, we will introduce some popular large margin algorithms and analyze their performance under our parallel framework.

3.1 MIRA

Crammer and Singer (2003) developed a large margin algorithm called MIRA and later extended by Taskar et.al (?). The algorithm has been widely used in many popular models (McDonald et al., 2005). It tries to minimize $\|w\|$ so that the margin between output score $s(x, z)$ and correct score $s(x, y)$ is larger than the loss of output structure:

$$\text{minimize} \|w\|$$

$$\text{st. } \forall z \in GEN(x) \ s(x, y) - s(x, z) \geq L(y, z)$$

During the inference process, the algorithm manages to find out the output structure with the highest score. However, in our parallel framework weight vector can be overwritten so we may not get the 1-best structure. Actually, it does not matter because the binary feature representation is so sparse that the output score is still close to 1-best score. We can say that the margin between output score and correct score is larger than the loss, so the margin between 1-best score and correct score will still satisfy the constrain.




Figure 1: Speed up of our parallel framework.

3.2 Structured Perceptron

Structured perceptron is first proposed by Collins (2002). It proves to be an effective and efficient structured prediction algorithm with convergence guarantee for separable data (Sun et al., 2009, 2013; Sun, 2015). We denote the binary feature representation of sequence x and structure y to be $f(x, y)$. The set of structure candidates for the input sequence x is denoted as $GEN(x)$. Structured perceptron searches the space of $GEN(x)$ and finds the output z with highest score $f(x, z) \cdot w$. The weight vector w then updates with the output z .

In our parallel framework, the inference and update of different samples runs parallel. The inference and update for structured perceptron can be described as:

$$z = \operatorname{argmax}_{t \in GEN(x)} f(x, t) \cdot w \quad (1)$$

$$w = w + f(x, y) - f(x, z) \quad (2)$$

During the inference, although the weight vector may be modified, we can still ensure that the score of the output z is higher than that of correct structure y . Huang et.al (2012) proved that if each update involves a violation (the output has a higher model score than the correct structure), structured perceptron algorithm is bound to converge. Therefore, our parallel framework on structure perceptron is effective theoretically.

4 Experiments

We compare our parallel framework and non-parallel large margin algorithms on several benchmark datasets.

4.1 Experiment Tasks

Part-of-Speech Tagging (POS-tag): Part of Speech Tagging is a famous and important task in

natural language processing. Following the prior work (Collins, 2002), we derive the dataset from Penn Wall Street Journal Treebank (Marcus et al., 1993). We use sections 0-18 of the treebank as training set while sections 19-21 is development set and sections 22-24 is test set. The selected feature is including unigrams and bigrams of neighboring words as well as lexical patterns of current word (Tsuruoka et al., 2011). We report the accuracy of output tag as evaluation metric.

Phrase Chunking (Chunking): In Phrase Chunking task, we tag the words in the sequence to be B, I or O to identify the noun phrases. The dataset is extracted from the CoNLL-2000 shallow-parsing shared task (Tjong Kim Sang and Buchholz, 2000). The feature includes word n-grams and part-of-speech n-grams. Our evaluation metric is F-score following prior works.

Biomedical Named Entity Recognition (Bio-NER): Biomedical Named Entity Recognition is mainly about the recognition of 5 kinds of biomedical named entities. The dataset is from MEDLINE biomedical text corpus. We use word pattern features and part-of speech features in our model (Tsuruoka et al., 2011). The evaluation metric is F-score.

4.2 Experiment Setting

We implement our parallel framework with large margin algorithms, including structured perceptron and MIRA on above benchmark datasets. We use development set to tune the learning rate α_0 and L2 regularization. The final learning rate α_0 is set as 0.02, 0.05, 0.005 for above three tasks, and L2 regularization is 1, 0.5, 5.

We implement a single-thread framework as our baseline. The setting of baseline is totally the same as our proposed framework. For fair com-

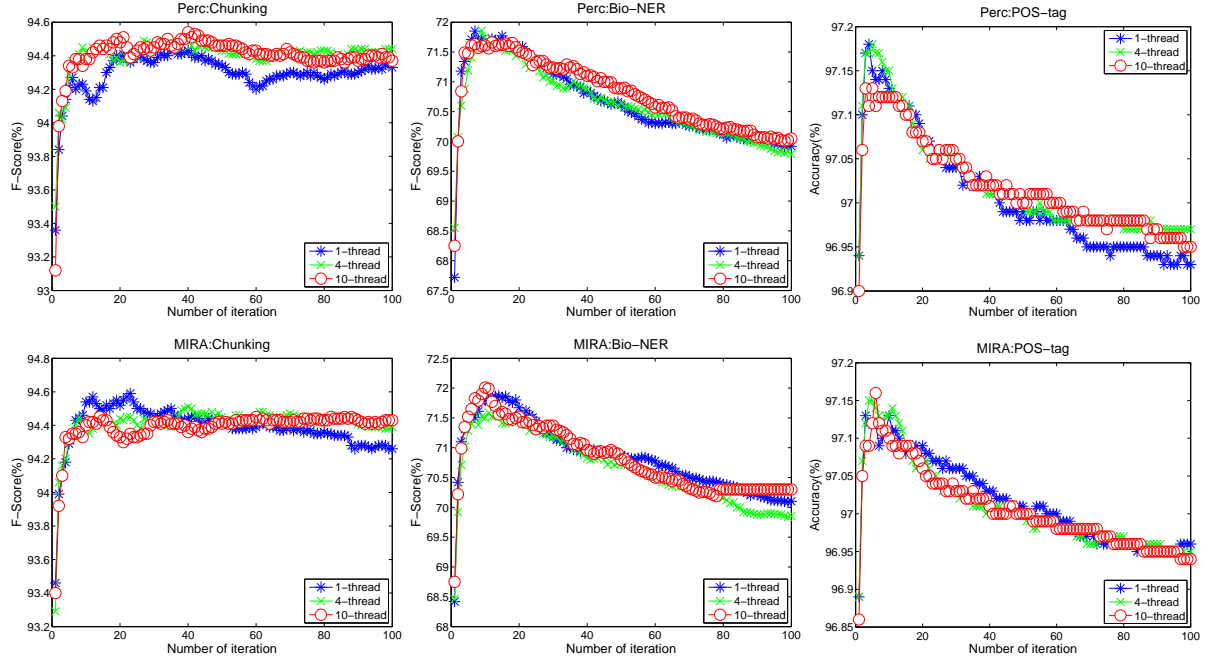


Figure 2: Experiment Results on the benchmark datasets.

Number of threads		1	4	10
Chunking	Perc	94.40	94.46	94.50
	MIRA	94.56	94.50	94.53
Bio-NER	Perc	71.83	71.90	71.80
	MIRA	71.75	71.70	71.91
POS-tag	Perc	97.17	97.18	97.10
	MIRA	97.13	97.15	97.15

Table 1: Accuracy/F-score of baseline and our framework.

Number of threads		1	4	10
Chunking	Perc	1.0x	3.0x	5.5x
	MIRA	1.0x	3.7x	4.7x
Bio-NER	Perc	1.0x	3.0x	5.0x
	MIRA	1.0x	3.0x	4.6x
POS-tag	Perc	1.0x	3.3x	4.4x
	MIRA	1.0x	3.4x	4.4x

Table 2: Speed up of our framework.

parison, we also average the parameters of the baseline. We run our parallel framework up to 10 threads following prior work (Recht et al., 2011). We compare our parallel framework with baseline in accuracy/F-score and time cost. Experiments are performed on a computer with Intel(R) Xeon(R) 3.0GHz CPU.

4.3 Experiment Result

Figure 1 shows that our parallel algorithm can gain near linear speed up. With 10 threads, our framework brings 4-fold to 6-fold faster speed than that with only 1 thread. Table 2 shows the speed up in our benchmark datasets with 1,4 and 10 threads.

Figure 2 also shows that our parallel framework has no loss in accuracy/F-score or convergence rate compared with single-thread baseline. Table 1 indicates that our framework does not hurt large margin algorithm because the difference of results is very small. In other words, there is barely interference among threads, and the strong robustness of large margin algorithm ensures no loss in performance under the parallel framework.

5 Conclusions

We propose a generic online parallel learning framework for large margin models. Our experiment concludes that the proposed framework has no loss in performance compared with baseline while the training speed up is near linear with increasing running threads.

6 Acknowledgements

This work was supported in part by National Natural Science Foundation of China (No. 61673028), and National High Technology Research and Development Program of China (863 Program, No.

2015AA015404). Xu Sun is the corresponding author of this paper.

References

- Michael Collins. 2002. Discriminative training methods for hidden markov models: Theory and experiments with perceptron algorithms. In *Proceedings of the ACL-02 conference on Empirical methods in natural language processing-Volume 10*. Association for Computational Linguistics, pages 1–8.
- Koby Crammer and Yoram Singer. 2003. Ultraconservative online algorithms for multiclass problems. *The Journal of Machine Learning Research* 3:951–991.
- Liang Huang, Suphan Fayong, and Yang Guo. 2012. Structured perceptron with inexact search. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, pages 142–151.
- Xiangru Lian, Yijun Huang, Yuncheng Li, and Ji Liu. 2015. Asynchronous parallel stochastic gradient for nonconvex optimization. In *Advances in Neural Information Processing Systems*. pages 2719–2727.
- Mitchell P Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of english: The penn treebank. *Computational linguistics* 19(2):313–330.
- Ryan McDonald, Koby Crammer, and Fernando Pereira. 2005. Online large-margin training of dependency parsers. In *Proceedings of the 43rd annual meeting on association for computational linguistics*. Association for Computational Linguistics, pages 91–98.
- Ryan McDonald, Keith Hall, and Gideon Mann. 2010. Distributed training strategies for the structured perceptron. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, pages 456–464.
- Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu. 2011. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems*. pages 693–701.
- Xu Sun. 2015. Towards shockingly easy structured classification: A search-based probabilistic online learning framework. *Technical report, arXiv:1503.08381*.
- Xu Sun. 2016. Asynchronous parallel learning for neural networks and structured models with dense features. In *COLING*.
- Xu Sun, Takuya Matsuzaki, and Wenjie Li. 2013. Latent structured perceptrons for large-scale learning with hidden information. *IEEE Trans. Knowl. Data Eng.* 25(9):2063–2075.
- Xu Sun, Takuya Matsuzaki, Daisuke Okanohara, and Jun’ichi Tsujii. 2009. Latent variable perceptron algorithm for structured classification. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*. pages 1236–1242.
- Erik F Tjong Kim Sang and Sabine Buchholz. 2000. Introduction to the conll-2000 shared task: Chunking. In *Proceedings of the 2nd workshop on Learning language in logic and the 4th conference on Computational natural language learning-Volume 7*. Association for Computational Linguistics, pages 127–132.
- Yoshimasa Tsuruoka, Yusuke Miyao, and Jun’ichi Kazama. 2011. Learning with lookahead: can history-based models rival globally optimized models? In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning*. Association for Computational Linguistics, pages 238–246.
- Shen-Yi Zhao and Wu-Jun Li. 2016. Fast asynchronous parallel stochastic gradient descent: A lock-free approach with convergence guarantee.
- Martin Zinkevich, Markus Weimer, Lihong Li, and Alex J Smola. 2010. Parallelized stochastic gradient descent. In *Advances in neural information processing systems*. pages 2595–2603.